

CLASSIFIED

AD 276 533

*Reproduced
by the*

VICES TECHNICAL INFORMATION AGENCY
ARLINGTON HALL STATION
ARLINGTON 12, VIRGINIA



CLASSIFIED

NOTICE: When government or other drawings, specifications or other data are used for any purpose other than in connection with a definitely related government procurement operation, the U. S. Government thereby incurs no responsibility, nor any obligation whatsoever; and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use or sell any patented invention that may in any way be related thereto.

CATALOGED BY ASTIA
AS AD NO.

276 533

276533

MEMORANDUM

RM-3173-PR

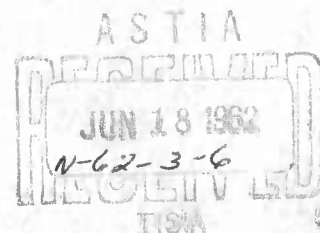
JUNE 1962

DYNAMIC PROGRAMMING,
INTELLIGENT MACHINES,
AND SELF-ORGANIZING SYSTEMS

Richard Bellman

PREPARED FOR:

UNITED STATES AIR FORCE PROJECT RAND



The **RAND** Corporation
SANTA MONICA • CALIFORNIA

MEMORANDUM

RM-3173-PR

JUNE 1962

DYNAMIC PROGRAMMING,
INTELLIGENT MACHINES,
AND SELF-ORGANIZING SYSTEMS

Richard Bellman

This research is sponsored by the United States Air Force under Project RAND — Contract No. AF 49(638)-700 — monitored by the Directorate of Development Planning, Deputy Chief of Staff, Research and Technology, Hq USAF. Views or conclusions contained in this Memorandum should not be interpreted as representing the official opinion or policy of the United States Air Force. Permission to quote from or reproduce portions of this Memorandum must be obtained from The RAND Corporation.

The **RAND** *Corporation*

1700 MAIN ST • SANTA MONICA • CALIFORNIA

PREFACE

In this Memorandum, the author discusses the importance of a scientific classification of problems in accordance with their structural features, if optimum use is to be made of intelligent machines for problem solving.

SUMMARY

In this Memorandum, the author discusses some aspects of problem formulation and problem solution^{*}. In particular, ~~he emphasizes~~ the relevance of these matters to the ~~exciting~~ field of intelligent machines and ~~indicates~~ some connections with the theory and application of dynamic programming^{**}. ~~Most important of all, he points out the~~ practical application of scientific philosophy as a technique to guide research and to avoid undue waste of time, energy, and talent. *is developed.*

** are discussed.*
*** are treated.*

CONTENTS

PREFACE111
SUMMARY	v
Section	
1. INTRODUCTION.	1
2. WHAT DO WE EXPECT FROM A THEORY?.	1
3. EXPENSIVE COMPARTMENTALIZATION.	2
4. GENERATION AND REGENERATION	3
5. HIERARCHIES OF PROCESSES.	4
6. SYSTEMS AND PROCESSES	4
7. INTELLIGENCE AND DECISION-MAKING.	5
8. FIRST-LEVEL PROCESSES	7
9. DISCUSSION.	8
10. INSTINCT.	8
11. SECOND-LEVEL PROCESSES.	9
12. INFINITE HIERARCHY OF PROCESSES	9
13. CREATIVITY.	10
14. INCONSISTENT LOGICS AND RANDOM AXIOM SYSTEMS. . .	11
15. MACHINES THAT PLAY CHESS IN THE NIGHT	11
16. DYNAMIC PROGRAMMING	13
17. MATHEMATICAL EXPERIMENTATION.	14
REFERENCES.	15

DYNAMIC PROGRAMMING, INTELLIGENT MACHINES AND SELF-ORGANIZING SYSTEMS

1. INTRODUCTION

As Gertrude Stein, author of such great truths as "A rose is a rose is a rose," lay on her deathbed, she was surrounded by acolytes and devotees who hoped to glean one last bit of wisdom from her. One of them, perhaps the faithful Alice B. Toklas, leaned near and asked, "Gertrude, Gertrude, what is the answer?" Summoning all of her remaining energy, Gertrude Stein opened her eyes and feebly responded, "What is the question?"

One cannot help feeling that too much effort has been devoted to obtaining answers without nearly enough effort being directed toward formulation of the proper question. Consequently, in what follows, we would like to discuss some aspects of problem formulation and problem solution. In particular, we shall emphasize the relevance of these matters to the exciting field of intelligent machines and indicate some connections with the theory and application of dynamic programming. Above all, we want to point out the practical application of scientific philosophy as a technique to guide research and to avoid undue waste of time, energy, and talent.

2. WHAT DO WE EXPECT FROM A THEORY?

A mathematician, like any intellectual, is not at all content with the mere collection of files of facts. He is well aware that the only defense against being overwhelmed by masses of data, accumulating alarmingly at an exponential rate each year, is an understanding of structural features, i.e., "theories," which enable us to organize vast domains of knowledge. Even this is not sufficient. As time goes on, theories themselves must be grouped in such a way as to recognize the basic identity

of numerous intellectual structures in different fields. As with nations, sciences have been kept needlessly apart by distinct languages.

If, then, we wish to depend upon theory for control and understanding of the physical universe, it is essential that we devote some effort to the intrinsic structure of theories, a theory of theories, as it were.

In constructing a theory, what do we demand in the way of problem solving? We ask the following:

- Recognition: The ability to discern and categorize questions which fall within the province of the theory.
- Formulation: The capacity for precise, but not necessarily equivalent, statements of the problem, using the vocabulary of the theory.
- Analysis: The possession of a set of systematic techniques for obtaining the structure of the solution in terms of the structure of the problem.
- Computation: The capability of furnishing numerical answers to numerical questions.

To all of these fundamental requirements, we wish to add one of operational significance:

- Regeneration: The ability to generate new questions from old answers; this implies flexibility, versatility and adaptability.

3. EXPENSIVE COMPARTMENTALIZATION

All of these characteristics are intimately intertwined, as we have emphasized in a number of papers [2], [3], [4]. Much of the difficulty that is experienced in applying mathematical methods to the problems of the physical world derives from the fact that there is seldom a unified attack upon a problem. Rather, the scientist—physical, social or biological—formulates on the basis of

his wide scientific knowledge, but limited mathematical skill; the mathematician analyzes on the base of broad mathematical skill, limited scientific background, and limited understanding of the abilities and limitations of digital computers. The computation is then done by people skilled in their specialty, but usually unaware of the physical or mathematical origins.

That anything should emerge from this sequence of poor impedance matching is remarkable. That the net effort is generally inefficient, inept, and irrelevant is not remarkable at all. With these matters in mind, it becomes clear that the dedication of time and effort to classification of problems is not an idle expenditure, but rather represents an important step in the direction of coordinating the work of many people in different fields. It is a particularly significant enterprise for those interested in intelligent machines, automata, adaptive processes, and decision theory.

Let us make one further point. We have tried in what follows to introduce some order in an area of great complexity, but we have not attempted any rigid axiomatic approach. There is conceivably some purpose in doing this in fields which have passed their heyday. In a living field, the aim of theory is to guide and stimulate, not to stultify and enchain.

4. GENERATION AND REGENERATION

Mathematics is one of the tools used by scientists to simplify and unify their domains. In return, mathematics is continually revitalized by the new problems arising from new scientific discoveries, e.g., quantum mechanics, relativity theory, digital computers. Without this continuing stimulus, mathematics withers and atrophies. It is not a self-sustaining field, although it can exist in embalmed and mummified form.

5. HIERARCHIES OF PROCESSES

Before turning to a classification of problems, and ipso facto, of solutions, let us point out that the idea we shall employ—that of a stratification of elements, a formation of hierarchies—is a standard one in mathematics.

We meet it in algebra and in recursive function theory. Indeed, a simple hierarchic argument shows the existence of transcendental numbers. In connection with the theory of sets, it can be used to establish the existence of continuous functions without derivatives at points, and so on.

The Liouville theory of elementary functions is based upon a sequential concept of operations. Finally, the theory of types was introduced by Russell (see the excellent expository article by Quine [5]), to avoid the kind of ambiguity that easily results from loose expression. We see similar meaninglessness today in discussions of "thinking machines."

6. SYSTEMS AND PROCESSES

We propose to categorize problems by associating them with various types of "systems" or "processes," terms whose significance will become clearer in context.

To describe a system we require some further concepts: state, cause and effect, and criterion, this latter only if we are talking about control processes. By the state of the system, we mean the parameters and functions which describe it. Thus, by means of simple concepts, a stone thrown into the air may be described at any time by its position and velocity. A more sophisticated and realistic theory will require further data. Hence, and a most important point, the state of a system is a relative concept—that is, relative to the mathematical and physical theories we are employing at the moment.

By cause and effect, we mean the change in the state of a system resulting from the application of a transformation. In the case of a control process, a decision is equivalent to a transformation, and we suppose that we know the new state resulting from the decision.

We wish to construct a hierarchy of processes and to indicate which ones we can handle in one way or another at the present time, and which ones are quite remote. In this way, we wish to introduce some precision into the concepts of "intelligence" and "thinking."

7. INTELLIGENCE AND DECISION-MAKING

We shall define intelligence as the ability to make decisions. It is thus easy to define levels of intelligence in terms of levels of decision processes. This interpretation of intelligence is reasonably consistent with that given by psychologists, namely, adaptation to one's environment.

We shall, for want of better, confine ourselves to decision processes involving numbers. The real difficulties in the applications of the quantitative methods of mathematics arise when we attempt to associate numbers with processes which do not intrinsically possess numerical utilities. We return to this point below.

It is important to distinguish clearly between level of computational or analytic complexity and conceptual or hierarchic level. There is conceptually no difficulty in evaluating 9 raised to the 99th power, but it would be difficult to do this quickly, and at the moment, we know no simple way of predicting the 723rd digit. Similarly, there is no conceptual difficulty in partial differential equations, although they may present severe analytic and computational difficulties. Frequently, the only way to overcome some of these obstacles is to use a higher conceptual level.

Thus, for example, to evaluate an integral such as

$$(7.1) \quad \int_{-\infty}^{\infty} \frac{\cos x dx}{1+x^2}$$

requires real ingenuity, but if we use the sophisticated tool of the theory of functions of a complex variable, it becomes a trivial exercise. Similarly, the "elementary" proof of the prime number theorem, i.e., without the help of complex variable theory, is quite difficult, whereas the usual proof is fairly straightforward.

In view of all of this, we may cite the theory of quasilinearization, where decision processes are artificially introduced to provide more efficient types of analytic approximation and computational solution. (See [6], [7], [8], and [10]. See also [9] where a two-person process is introduced to simplify the analysis.)

Nonetheless, there is value in constructing a hierarchy of processes, following very much the lines of the Liouville and Russell theories mentioned above.

At the simplest level, we have descriptive processes, characterized by the iteration of a fixed transformation

$$(7.2) \quad p_{n+1} = T(p_n), \quad p_0 = c.$$

There are in many cases formidable analytic difficulties in the path of carrying out these instructions, but there is no conceptual difficulty. Thus, as mentioned above, the calculation of 9^{99} requires a good deal of effort if we employ only simple arithmetic. Presumably, with a sufficient knowledge of number theory we can obtain more efficient techniques. We can thus regard one of the principal objectives of mathematical sophistication as the overcoming of arithmetic. Anytime we are reduced to brute calculation, we have made an admission of weakness.

Descriptive processes of the foregoing type will be called zero-level processes.

8. FIRST-LEVEL PROCESSES

The first significant step in the direction of intelligent machines was made in the application of the concept of feedback control. Although an ancient idea, it was first fully brought into prominence by the Industrial Revolution. The many intriguing mathematical aspects of this concept were early recognized by Maxwell [11] and Vishnegradsky [12]; see [13] for a reprint of Maxwell's paper and other fundamental papers in control theory.

Referring to the analytic representation of (7.2), we now assume that at each stage of the process there exists a set of available transformations, $T(p,q)$, where q designates the individual transformation. "Feedback" means that the transformation which is applied depends in some fashion upon the current state of the system, i.e., $q = q(p)$.

We can thus write as the analytic equivalent of a process of the first level,

$$(8.1) \quad p_{n+1} = T(p_n, q(p_n)).$$

In many cases, this choice of a transformation (or "policy") is determined by an optimization criterion. In most cases, we can pretend that such an optimization process exists and interpret the actual policy in these terms. This is the standard "as if" technique of mathematical physics exemplified by Fermat's Principle, Hamilton's Principle, etc.

Similarly, in the study of biological systems we may pretend that organisms behave as if they were trying to maximize their probability of survival.

9. DISCUSSION

If we consider only deterministic processes, we find that first-level processes collapse to zero-level processes, since

$$(9.1) \quad p_{n+1} = T(p_n, q(p_n)) = T_1(p_n),$$

again a transformation of fixed type. Conversely, we can in many cases gain considerable advantage from writing (7.2) in the form of (9.1) above; see [6], [7], and the discussion in Sec. 7.

This reduction in hierarchic rank is due to the duality between point and line formulation of Euclidean geometry, an equivalence which masks the new idea of feedback control; see [1] for a further discussion.

Only when we turn to stochastic decision processes do we see the essential difference between level zero and level one, and recognize the true meaning of "feedback." This meaning was obscured for many years.

10. INSTINCT

Much of the classical argument concerning the difference between "Instinct" and "Intelligence" can be resolved by definition, using the foregoing concepts. What is generally called instinct can be considered to be intelligence of level one, as defined above.

Similarly, a number of choice philosophical and theological conundrums such as "Free Will" versus "Predestination" can be seen to be unanswerable because ill-defined and ill-posed questions have been asked. Naturally, definitions never completely resolve fundamental questions, but a formulation of the type we are presenting enables us to pinpoint basic issues with more accuracy.

The problem of instinct is clearly allied with that of information pattern and discrimination. When the

organism cannot distinguish between real and spurious stimuli, we see such effects as the mass suicide of the lemmings, and allergic and neurotic behavior.

11. SECOND-LEVEL PROCESSES

We have defined zero-level processes, and then first-level processes as those requiring a decision as to the operation of the zero-level process. The stage has then been set for an inductive or recursive definition of higher-level processes.

A second-level process will be defined as one requiring a decision concerning the operation of a first-level process. We encounter processes of this type in connection with simple learning and in adaptive control theory.

Consider a stochastic control process of the following type. At each stage there is a random effect which tends to disturb the system. To correct for this deviation, we use feedback control. Suppose that the random effect is represented by a random variable r which has the simple distribution

$$r = +1, \text{ with probability } p,$$

$$r = -1, \text{ with probability } 1 - p.$$

If p is known, we have a process of level one. If p is fixed, but unknown, we have a process of level two. Here, on the basis of the history of the process over time, we make successively better estimates of p .

A detailed discussion of processes of this nature will be found in [1].

12. INFINITE HIERARCHY OF PROCESSES

If we use the preceding formulation, it is now easy to construct a hierarchy of processes, each requiring a more

sophisticated adaptation. Consider the case where p , the unknown probability, is known to have a distribution function $dG(p,a)$. Here a is an unknown parameter, which may itself (to introduce a higher level process) be considered to have a distribution function $dH(a,b)$, which depends upon an unknown parameter b , and so on.

To see how situations of this type arise in practice, suppose that we are picking coins of the same kind from a barrel, where each coin in the barrel has the same probability of coming up heads. Now suppose that there is a room full of barrels of coins, each barrel having its own probability distribution, and a building full of barrel-packed rooms, a warehouse center full of such buildings, and so on.

We see then that we can construct a hierarchy of decision processes and thus a hierarchy of levels of intelligence. Simultaneously, using standard procedures, we can construct decision processes not belonging to any member in this hierarchy. Where, for example, do we put the problem of determining at what level a particular process belongs?

13. CREATIVITY

It appears that if we wish to talk about creativity in some reasonable fashion, we must introduce similar concepts to those mentioned above. We want to distinguish between innovation within the framework of a particular decision level and innovation outside of this framework. It is not difficult to see that we possess very feeble means, if any, for transcending hierarchies and that to contemplate this type of activity by means of digital computers is rather absurd. The use of random stimuli to simulate human thinking is too reminiscent of Gulliver's experiences in the laboratories of Laputa to be taken seriously.

14. INCONSISTENT LOGICS AND RANDOM AXIOM SYSTEMS

Classical logical systems have insisted upon the development of approaches to problem solving which yield correct answers without fail. This is a luxury to which one cannot aspire in more complex situations, principally because we do not understand enough about them to formulate the problems in completely meaningful terms. We consequently develop a set of local logics each applicable to a neighborhood of processes. These need not be consistent when they overlap, and in general will not be. There is thus nothing more illogical than trying to carry arguments to their "logical" conclusion in the political, social and economic spheres. It is a denial of the history of science to look for "unified" theories of complex phenomena. If it cannot be done in the physical sciences, how then in the more difficult fields of human behavior?

This does not mean that we allow irrationality freely. On the contrary, the whole endeavor of the intellectual approach, the so-called scientific method, is to reduce the regions of irrationality. If we are interested in actually solving problems by means of digital computers, it may be more efficient to use a number of different approximate logics, with constant control and comparison. If we want to minimize the time required to resolve a question, we may be able to accomplish this with a set of approximate policies, tried one after the other, rather than with one ponderous technique which has complete certitude.

Closely allied with this is the investigation of situations in which a variety of consistent logical systems are used in a random fashion. Simulation processes of this nature could produce some interesting results.

15. MACHINES THAT PLAY CHESS IN THE NIGHT

The best chess-playing machines to date were built in the nineteenth century. They used the very simple idea of

having a human hidden inside. With the introduction of large-scale computers, the idea of constructing electronic chess players was pursued again by a number of enthusiasts. This is a very intriguing idea, but—like many ideas of this nature—one that can lead to a substantial waste of time and effort if no adequate preliminary evaluation is made.

It is rather unfortunate that digital computers, which were designed by such mathematical pioneers as Wiener and von Neumann to resolve a number of significant scientific questions, have to such a great extent been relegated to accounting and bookkeeping and to the routine solution of routine engineering problems. It is even more unfortunate that the small amount of time left over from these profitable (?) activities has been devoted to the calculation of prime numbers and to the development of chess-playing programs.

Why is chess a poor game to analyze by computer?

In the first place, it is not a game, such as tic-tac-toe, which can be resolved by merely an enumeration of possible combinations of moves. Even the game of checkers, far more limited in strategic and tactical concepts, cannot be played on a master level by computer, although sophisticated enumerative and feedback techniques can be used to program a computer to play a strong game.

Chess requires a deeper understanding—which we do not possess at the present time. Essentially the difficulty is that tactics and strategy intertwine in complex ways. Although the openings are fairly standard, the middle game is an uncharted jungle.

In natural situations in which we do not possess precise optimal policies, we must employ approximate techniques. This is where the unnatural character of chess as a contrived process manifests itself. Chess is an inherently unstable game. By this we mean that a small difference in the

position of the pieces and pawns can make the difference between winning and losing. Even the number of pieces and pawns is frequently of no significance. Thus approximation techniques are of questionable value.

The concept of checkmate is an interesting but highly artificial one. Just the opposite is true of physical processes where there is necessarily stability, a fundamental observation of Hadamard, and where approximate policies work extremely well—fortunately for physicists and engineers.

Let us note, however, that this opinion of the feasibility of constructing chess-playing programs is not shared by the world champion, Botwinnik, who is actually engaged in investigations of this type. On the other hand, another great master, Edward Lasker, shares the opinions expressed above.

16. DYNAMIC PROGRAMMING

As we have already mentioned, the hierarchy of decision processes we constructed was considerably influenced by the a priori knowledge of a mathematical theory that could formulate the resultant mathematical problems in precise analytic terms, and, occasionally, in happy circumstances, resolve them analytically and computationally.

The basic idea is that of a policy which is a sequence of decisions, each of which is based upon current information. This information pattern, or state variable, is the clue to the complexity of the decision process. If the state variable is an n -dimensional vector, or even a function, we have a process of level one. If it is a distribution function for a vector in n -dimensional space, we have a process of level two. If it is a distribution function for parameters which determine a distribution function, we have a process of level three and so on.

We see then that we could employ the Liouville type of classification to construct a hierarchy, without any reference to the original decision processes. This is convenient, but an identification must still be made at some point.

Detailed discussion of deterministic, stochastic, and adaptive processes will be found in [1].

17. MATHEMATICAL EXPERIMENTATION

Since we know so little about self-organizing systems, we must carry out a large number of trial investigations. In other words, we must engage in mathematical experimentation. One of the most useful tools the mathematician possesses is the digital computer.

At the present time, a good deal of work is going on in connection with learning and adaptive processes, but not nearly enough compared to the magnitude of the difficulties and the great variety of different problems confronting us.

This does not mean that we should search randomly. It does mean that we should investigate the behavior of specific mathematical models, the efficiency of various types of approximation, particularly approximation in policy space, and so on.

In particular, we must find simple approximate techniques for reducing hierarchic level and for using as much of a quantitative method as possible in the study of qualitative phenomena. If we are to avoid the morass of metaphysics, we must reduce as many concepts as possible to numerical terms. On the other hand, we must face the fact that the most important aspects of human life are intrinsically nonnumerical. Any attempt to ignore this is highly unscientific. In the true intellectual approach, one accepts this fact and copes with it.

REFERENCES

1. Bellman, R., Adaptive Control Processes: A Guided Tour, The RAND Corporation, R-350, 1961; also published by Princeton University Press, Princeton, New Jersey, 1961.
2. Bellman, R., "Mathematical Experimentation and Biological Research," Federation Proceedings, Vol. 21, 1962, pp. 109-111.
3. Bellman, R., and P. Brock, "On the Concepts of a Problem and Problem-solving," Amer. Math. Monthly, Vol. 67, 1960, pp. 119-134.
4. Bellman, R., C. Clark, C. Craft, D. Malcolm, and F. Ricciardi, "On the Construction of a Multi-person, Multi-stage Business Game," Operations Research, Vol. 5, 1957, pp. 469-503.
5. Quine, W. V., "Paradox," Scientific American, Vol. 206, April 1962, pp. 84-100.
6. Bellman, R., "Functional Equations in the Theory of Dynamic Programming—V: Positivity and Quasi-linearity," Proc. Nat. Acad. Sci. USA, Vol. 41, 1955, pp. 743-746.
7. Kalaba, R., "On Nonlinear Differential Equations, the Maximum Operation, and Monotone Convergence," J. Math. and Mech., Vol. 8, 1959, pp. 519-574.
8. Kalaba, R., Some Aspects of Quasilinearization, The RAND Corporation, RM-2924-PR, December 1961.
9. Bellman, R., I. Glicksberg, and O. Gross, "Some Non-classical Problems in the Calculus of Variations," Proc. Amer. Math. Soc., Vol. 7, 1956, pp. 87-94.
10. Bellman, R., "Quasilinearization and Upper and Lower Bounds for Variational Problems," Q. Appl. Math., Vol. 19, 1962, pp. 349-350.
11. Maxwell, J. C., "On Governors," Proc. Royal Soc. London, Vol. 16, 1868, pp. 270-283.
12. Vishnegradsky, J., "Sur la theorie generale des Regulateurs," Compt. Rend. Acad. Sci. Paris, Vol. 83, 1876, pp. 318-321.
13. Bellman, R., and R. Kalaba, Mathematical Trends in Control Theory, Dover Publications, New York, to appear in 1962.